UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/017,371	12/07/2001	Leith Johnson	10016615-1	8105

22879 7590 01/07/2008

HEWLETT PACKARD COMPANY
P O BOX 272400, 3404 E. HARMONY ROAD
INTELLECTUAL PROPERTY ADMINISTRATION
FORT COLLINS, CO 80527-2400

EXAMINER	
TSAI, SHENG JEN	
ART UNIT	**PAPER NUMBER**
2186	

NOTIFICATION DATE	DELIVERY MODE
01/07/2008	ELECTRONIC

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Notice of the Office communication was sent electronically on above-indicated "Notification Date" to the following e-mail address(es):

JERRY.SHORMA@HP.COM
mkraft@hp.com
ipa.mail@hp.com

UNITED STATES PATENT AND TRADEMARK OFFICE

Commissioner for Patents
United States Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450
www.uspto.gov

BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

Application Number: 10/017,371
Filing Date: December 07, 2001
Appellant(s): JOHNSON, LEITH

MAILED
JAN 0 3 2008
Technology Center 2100

Robert Plotkin
For Appellant

EXAMINER'S ANSWER

This is in response to the appeal brief filed 10/17/2007 appealing from the Office action mailed 3/5/2007.

(1) Real Party in Interest

A statement identifying by name the real party in interest is contained in the brief.

(2) Related Appeals and Interferences

The examiner is not aware of any related appeals, interferences, or judicial proceedings which will directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal.

(3) Status of Claims

The statement of the status of claims contained in the brief is correct.

(4) Status of Amendments After Final

The appellant's statement of the status of amendments after final rejection contained in the brief is correct.

(5) Summary of Claimed Subject Mater

The summary of claimed subject matter contained in the brief is correct.

(6) Grounds of Rejection to be Reviewed on Appeal

The appellant's statement of the grounds of rejection to be reviewed on appeal Identifies the ground of rejections and the associated claims under rejection to be reviewed on appeal.

(7) Claims Appendix

The copy of the appealed claims contained in the Appendix to the brief is correct.

(8) Evidence Relied Upon

5,860,146	Vishin et al.	01-1999

(9) Grounds of Rejection

Claims 1-5, 8-11, 14-21, 23, 25 and 27-31 are rejected under 35 U.S.C. 102(b) as being anticipated by Vishin et al. (US 5,860,146).

As to claim 1, Vishin et al. disclose **in a partitionable computer system** [figure 1 shows a partitionable computer system: the computer system is partitioned into a plurality of clusters (102), and each cluster is further partitioned into a plurality of processing unit (104)] **including a plurality of machine resources having a plurality of machine resource identifiers** [figure 1 shows a partitionable computer system: the computer system is partitioned into a plurality of clusters (102), and each cluster is further partitioned into a plurality of processing unit (104); each of the plurality of clusters represent a plurality of machine resources (i.e., nodes) with machine resource identifier (the Node-ID, figure 7, 170)], **a method for creating a physical resource identifier space in a partition of the partitionable computer system** [the corresponding physical resource identifier space is the Remote Translation Lookaside Table (RTLB, figure 8, 160) which is used to derive the Remote Page Physical Address (figure 6, 168; figure 7) that is then used to access Remote Physical Address (figure 5) associated with the plurality of clusters (figure 1, 102); abstract], **the method comprising steps of**:

(A) **establishing a mapping** [figures 2-8 show the mapping; abstract] **between a plurality of physical resource identifiers** [the RTLB, figure 8, 160] **and at least some of the plurality of machine resource identifiers** [Node-ID (figure 7, 170) and the associated Remote Physical Address (figure 7, 172)], **wherein the plurality of**

physical resource identifiers are numbered sequentially beginning with zero [figure 8, 160 shows that the RTLB includes a plurality of RPTE entries that are numbered sequentially beginning with 0]; **and wherein the mapping defines a non-monotonic function** [figure 8 shows that the mapping from the RPTE to the Remote Physical Address is non-monotonic, because the RPTE of entry #1 is mapped into the slashed area (the file and DB segment locked by local processor) located in the middle of the Remote Physical Address Space while the RPTE of entry #8 is mapped into areas of the Remote Physical Address Space both "above" and "below" that of the RPTE of entry #1. Thus this mapping is clearly non-monotonic]; **and**

(B) **providing, to an operating system** [the operating system, figure 9, 180] **executing in the partition** [In response to the page fault the operating system of the cluster 102 will request the memory controller 112 to retrieve the specified page from secondary memory 110 and store it in a free page in primary memory 108 (column 2, lines 36-40)], **an interface for the operating system to access the at least some of the plurality of machine resources using the plurality of physical resource identifiers as indices into the content address memory** [figure 9 shows the interface configuration for the operating system (180) to access the RTLB (160) within the memory controller (112), to access the TLB (122) within the Proc (104) via the cache (120), and to access other clusters (102) via the communication interface unit and the network interconnectivity unit (114); The 32 entries in the RTLB are organized into four groups of eight entries (entries 0:7, 8:15, 16:23 and 24:31). While the physical address ranges of the RPTEs in any one group may overlap, it is the responsibility of

the operating system 180 (see FIG. 9) to make sure that the RPTEs in different groups do not have overlapping address ranges (column 5, lines 33-59)].

As to claim 2, Vishin et al. disclose that **the plurality of machine resources comprises a plurality of machine memory locations, wherein the plurality of machine resource identifiers comprises a plurality of machine memory addresses** [as shown in figures 1-8; figure 1 shows a partitionable computer system: the computer system is partitioned into a plurality of clusters (102), and each cluster is further partitioned into a plurality of processing unit (104); each of the plurality of clusters represent a plurality of machine resources (i.e., nodes) with machine resource identifier (the Node-ID, figure 7, 170)], **and wherein the plurality of physical resource identifiers comprises a plurality of physical memory addresses** [figure 7 shows Node-ID (figure 7, 170) and the associated Remote Physical Address (figure 7, 172); figures 1-8 illustrate the mapping into the physical address associated with the plurality of machine resources (the clusters, figure 1, 102)].

As to claim 3, Vishin et al. disclose that **the method of claim 1 further comprising a step of performing the steps (A) and (B) for each of a plurality of partitions of the partitionable computer** [figure 7 shows Node-ID (figure 7, 170) and the associated Remote Physical Address (figure 7, 172); figures 1-8 illustrate the mapping into the physical address associated with the plurality of machine resources (the clusters, figure 1, 102)].

As to claim 4, Vishin et al. disclose that **the step (A) comprises a step of creating an address translation table that records the mapping between the**

plurality of physical resource identifiers and the at least some of the plurality of machine resource identifiers [the Remote Translation Lookaside Table (RTLB, figure 8, 160) which is used to derive the Remote Page Physical Address (figure 6, 168; figure 7) that is then used to access Remote Physical Address (figure 5) associated with the plurality of clusters (figure 1, 102); abstract].

As to claim 5, Vishin et al. disclose that **the interface comprises means for translating a physical resource identifier selected from among the plurality of physical resource identifiers into one of the plurality of machine resource identifiers in accordance with the mapping** [the Remote Translation Lookaside Table (RTLB, figure 8, 160) which is used to derive the Remote Page Physical Address (figure 6, 168; figure 7) that is then used to access Remote Physical Address (figure 5) associated with the plurality of clusters (figure 1, 102); abstract].

As to claim 8, refer to "As to claim 1."

As to claim 9, refer to "As to claim 2."

As to claim 10, refer to "As to claim 4."

As to claim 11, refer to "As to claim 5."

As to claim 14, refer to "As to claim 1" through "As to claim 5." Further, Vishin et al. disclose that [Each remote page table entry represents a mapping between a range of physical addresses and a corresponding range of remote physical addresses. The primary translation lookaside buffer translates a virtual address asserted by the data processor into a physical address. When the physical address does not correspond to a location in the local memory, the RTLB determines whether the physical address

matches at least one of the remote page table entries stored in the RTLB, and selects one of those remote page table entries when at least one match is found. The RTLB's selection circuitry selects a single remote page table entry in accordance with predefined RPTE selection criteria when two or more of the remote page table entries match the physical address. Then, a remote physical address is generated by combining a portion of the selected remote page table entry with a portion of the physical address (abstract); The present invention relates generally to multiprocessor computer systems having virtual memory management subsystems, and particularly to a memory controller that manages access to remote physical addresses through the use of an auxiliary translation lookaside buffer (column 1, lines 5-10)].

As to claim 15, refer to "As to claim 2."

As to claims 16-17, Vishin et al. disclose that [For instance, the control flags 176 may be used to indicate "read only" access or "read/write" access to the specified remote memory block (column 7, lines 48-53)].

As to claim 18, refer to "As to claim 14."

As to claim 19, refer to "As to claim 2."

As to claims 20-21, refer to "As to claims 16-17."

As to claim 23, refer to "As to claim 1" through "As to claim 5." Further, Vishin et al. disclose

(A) selecting a first subset of the plurality of physical memory locations, the first subset of the plurality of memory locations being mapped to a first subset of the plurality of machine memory addresses [When the physical address does not

correspond to a location in the local memory, the RTLB determines whether the physical address matches at least one of the remote page table entries stored in the RTLB, and selects one of those remote page table entries when at least one match is found. The RTLB's selection circuitry selects a single remote page table entry in accordance with predefined RPTE selection criteria when two or more of the remote page table entries match the physical address. Then, a remote physical address is generated by combining a portion of the selected remote page table entry with a portion of the physical address (abstract)]; **and**

(B) copying the contents of the first subset of the plurality of machine memory addresses to the second subset of the plurality of machine memory addresses [If the cluster's memory 108 does not store the page of the remote page table containing the required RPTE, that page of the remote page table 150 will first need to be downloaded from an appropriate remotely located processor or cluster in the distributed system 100. Once the required RPTE 152 is found and the address translation into a remote physical address (RPA) is performed, then a request is transmitted via the network 114 to load a copy of the page being addressed into the requesting cluster's local memory 108 (column 3, lines 22-30)].

As to claim 25, refer to "As to claim 23."

As to claim 27, Vishin et al. teach that **the interface comprises a Content Addressable Memory that establishes the mapping** [figure 6, 162 shows that a CAM is used to implement a remote translation lookaside buffer (RTLB); column 4, lines 63-67; column 5, lines 1-20, and the association of the physical resource

identifiers and the indices of the CAM (figure 6; column 5, lines 60-67; column 4, lines 5-10; column 5, lines 32-59; column 6, lines 1-35)].

As to claim 28, refer to "As to claim 27."

As to claim 29, refer to "As to claim 27."

As to claim 30, Vishin et al. teach **copying the contents of the first subset of the plurality of machine memory addresses to the second subset of the plurality of machine memory addresses** [If the cluster's memory 108 does not store the page of the remote page table containing the required RPTE, that page of the remote page table 150 will first need to be downloaded from an appropriate remotely located processor or cluster in the distributed system 100. Once the required RPTE 152 is found and the address translation into a remote physical address (RPA) is performed, then a request is transmitted via the network 114 to load a copy of the page being addressed into the requesting cluster's local memory 108 (column 3, lines 22-30)].

As to claim 31, refer to "As to claim 30."

(10) Response to Arguments

Appellants' arguments have been fully and carefully considered with Examiner's answers set forth below.

Response to Argument on Claims 1-5, 8-11, 14-21 and 28-29

Appellants contend that the Vishin reference fails to teach the limitation of "providing, to an operating system executing in the partition, an interface for the operating system to access at least some of the plurality of machine resources using the plurality of physical resource identifiers," because Vishin does not disclose any

"interface" that is provided to the operating system to do so, and that Vishin does not disclose that the operating system accesses machine resources. The Examiner disagrees with this argument for the following reasons:

First, according to the definition provided by Microsoft Computer Dictionary (Microsoft Press, 5th edition, 2002, page 279), an "interface" is "the point at which a connection is made between two elements so that they can work with each other or exchange information." Vishin shows in figure 9 the connections for the operating system (180) to access the RTLB (160) within the memory controller (112), to access the TLB (122) within the Proc (104) via the cache (120), and to access other clusters (102) via the communication interface unit and the network interconnectivity unit (114).

Thus, Vishin clearly discloses in figure 9 an "interface" for the operating system (figure 9, 180) to access all the machine resources that are illustrated in figures 1-8, including the RTLB (figure 1, 160; figure 6, 160; figure 8, 160; figure 9, 160), the TLB (figure 1, 122; figure 9, 122), the DRAM (figure 1, 108; figure 2, 108; figure 4, 108), the SRAM (figure 6, 166) and the clusters (figure 1, 102; figure 9, 102).

Second, "accessing" a resource does not have to be done in a "directly connected" manner. For example, when a user "accesses" various web sites through the Internet the user is only directly connected to the keyboard of the computer being used, and the user is not directly connected to the content resources of those web sites. The "target" to be accessed is the content resources of those web sites, and all the links between the user and the content resources of those web sites constitutes the "interface" connecting between the user and the web sites.

Thus, when the operating system (figure 9, 180) requests the memory controller (figure9, 112) to retrieve specific pages from secondary memory (figure 1, 110), as disclosed by Vishin [In response to the page fault the operating system of the cluster 102 will request the memory controller 112 to retrieve the specified page from secondary memory 110 and store it in a free page in primary memory 108 (column 2, lines 36-40); While the physical address ranges of the RPTEs in any one group may overlap, it is the responsibility of the operating system 180 (see FIG. 9) to make sure that the RPTEs in different groups do not have overlapping address ranges (column 5, lines 33-59)], The "target" to be accessed by the operating system is the specific pages from secondary memory (figure 1, 110), and the memory controller (figure9, 112) and all the links shown in figure 9 between the operating system (figure 9, 180) and the specific pages from secondary memory constitutes the "interface" connecting between the operating system and the secondary memory.

Thus, these passages provide concrete examples as to how the operating system utilizes the "interface" configuration illustrated in figure 9 to reach the various machine resources.

Appellants further cite, by referring to certain passages of the specification, examples of "the interface" as supporting evidences that Vishin fails to teach the element "an interface" recited in claim 1.

However, the claim language of claim 1 merely recites "an interface," and is otherwise completely silent regarding what constitutes "an interface" or any specific characteristics that this "interface" must possess. As such, the Examiner has to give the

element "interface" its broadest possible and reasonable interpretation, as required by MPEP. The definition of "interface," cited above and being applied by the Examiner for the purpose of claim analysis with respect to prior art, represents the broadest possible and reasonable interpretation of the term.

Appellants are reminded that only those limitations explicitly and expressly recited in a claim are considered by the Examiner during claim analysis. The written description of the specification may cite numerous examples to provide further information on the invention. While these numerous examples may prove to be informative in helping persons of ordinary skills in the art to better understand the invention, they are not taken into consideration when determining the patentability of a claim because they are not recited in a claim and are not part of the limitations of a claim.

Third, as to the allegation that Vishin does not teach that the operating system accesses machine resources, it is note that Vishin teaches that [In response to the page fault the operating system of the cluster 102 will request the memory controller 112 to retrieve the specified page from secondary memory 110 and store it in a free page in primary memory 108 (column 2, lines 36-40)], and that [While the physical address ranges of the RPTEs in any one group may overlap, it is the responsibility of the operating system 180 (see FIG. 9) to make sure that the RPTEs in different groups do not have overlapping address ranges (column 5, lines 33-59)].

Appellants contend that the operating system (figure 9, 180) only access the memory controller (figure 9, 112) and does not access memory using a plurality of

physical resource identifiers. Note that Vishin teaches "In response to the page fault the operating system of the cluster 102 will request the memory controller 112 to retrieve the specified page from secondary memory 110 and store it in a free page in primary memory 108" (column 2, lines 36-40). Since the "memory controller" is part of the "interface" as shown in figure 9, the fact that "The operating system requests the memory controller to retrieve the specified page from the secondary memory" is consistent with the Examiner's interpretation that "the operating system" uses the "interface" (i.e., the memory controller) to access the specified page from the secondary memory. Note that the specified page corresponds to the limitation "physical resource identifiers."

Thus Vishin clearly illustrates by the above examples how the operating system can access the machine resources.

Therefore, the Examiner's position regarding the patentability of these claims remains the same as stated in the previous Office Action.

(11) Related Proceedings Appendix

None.

Sheng-Jen Tsai Examiner Art Unit 2186 Sheng-Jen Tsai
/Matthew M. Kim/ Matthew Kim Supervisory Patent Examiner Art Unit 2186

/Lynne H Browne/
Lynne H Browne
Appeal Practice Specialist, TQAS
Technology Center 2100

December 20, 2007